

An Analysis of Technical Observations in Insider Theft of Intellectual Property Cases


Michael Hanley
Tyler Dean
Will Schroeder
Matt Houy
Randall F. Trzeciak
Joji Montelibano

February 2011

TECHNICAL NOTE
CMU/SEI-2011-TN-006

CERT® Program
Unlimited distribution subject to the copyright.

<http://www.sei.cmu.edu>



This report was prepared for the

SEI Administrative Agent
ESC/XPK
5 Eglin Street
Hanscom AFB, MA 01731-2100

The ideas and findings in this report should not be construed as an official DoD position. It is published in the interest of scientific and technical information exchange.

This work is sponsored by the U.S. Department of Defense. The Software Engineering Institute is a federally funded research and development center sponsored by the U.S. Department of Defense.

Copyright 2011 Carnegie Mellon University.

NO WARRANTY

THIS CARNEGIE MELLON UNIVERSITY AND SOFTWARE ENGINEERING INSTITUTE MATERIAL IS FURNISHED ON AN "AS-IS" BASIS. CARNEGIE MELLON UNIVERSITY MAKES NO WARRANTIES OF ANY KIND, EITHER EXPRESSED OR IMPLIED, AS TO ANY MATTER INCLUDING, BUT NOT LIMITED TO, WARRANTY OF FITNESS FOR PURPOSE OR MERCHANTABILITY, EXCLUSIVITY, OR RESULTS OBTAINED FROM USE OF THE MATERIAL. CARNEGIE MELLON UNIVERSITY DOES NOT MAKE ANY WARRANTY OF ANY KIND WITH RESPECT TO FREEDOM FROM PATENT, TRADEMARK, OR COPYRIGHT INFRINGEMENT.

Use of any trademarks in this report is not intended in any way to infringe on the rights of the trademark holder.

Internal use. Permission to reproduce this document and to prepare derivative works from this document for internal use is granted, provided the copyright and "No Warranty" statements are included with all reproductions and derivative works.

External use. This document may be reproduced in its entirety, without modification, and freely distributed in written or electronic form without requesting formal permission. Permission is required for any other external and/or commercial use. Requests for permission should be directed to the Software Engineering Institute at permission@sei.cmu.edu.

This work was created in the performance of Federal Government Contract Number FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. The Government of the United States has a royalty-free government-purpose license to use, duplicate, or disclose the work, in whole or in part and in any manner, and to have or permit others to do so, for government purposes pursuant to the copyright license under the clause at 252.227-7013.

For information about SEI publications, please visit the library on the SEI website (www.sei.cmu.edu/library).

Table of Contents

Acknowledgments	vii
Executive Summary	ix
Abstract	xi
1 Introduction	1
1.1 Background	1
1.2 Definitions	1
1.3 CERT® Insider Threat Database	2
1.4 Considerations	4
2 Analysis	6
2.1 Case Pool	6
2.2 Nontechnical Findings of Interest	6
2.3 Data Structure	6
2.4 Data Analysis by Field	7
2.4.1 Case Name and ID	7
2.4.2 Case Summary	7
2.4.3 Assets Attacked/Targeted	7
2.4.4 Methods of Exfiltration	9
2.5 Analysis Across Multiple Fields	13
2.5.1 Concealment	15
3 Considerations for Mitigation	16
3.1 Recommendations	16
4 Conclusion and Future Work	18
References	19

List of Figures

Figure 1:	Number of Cases in the CERT Insider Threat Databases by High-Level Category (Excluding National Security Espionage Cases)	3
Figure 2:	Cases in Three Major Crime Types by Sector	4
Figure 3:	Assets Attacked by Insiders	9
Figure 4:	Exfiltration Methods (Aggregate)	10
Figure 5:	Exfiltration by Asset Type	14

List of Tables

Table 1:	Types of Insider Crimes	2
Table 2:	Sample Case Table	7
Table 3:	Assets Attacked/Targeted Responses	8
Table 4:	Exfiltration Methods	10

Acknowledgments

Special thanks to Carnegie Mellon CyLab for partially funding this work. Also, thanks to Andrew P. Moore, Craig Lewis, Dawn Cappelli, and Paul Ruggiero, all of the Software Engineering Institute's CERT Program.

Executive Summary

The insider threat problem has recently been garnering more interest in business and government. Employees can threaten an organization's control of national secrets, trade secrets, and intellectual property (IP). The Insider Threat database at the CERT[®] Program, part of Carnegie Mellon University's Software Engineering Institute, catalogs more than 550 actual insider cases, including theft of IP, IT sabotage, and fraud. These case studies can provide greater insight into trends, detailed events, actions, and conditions in actual insider crimes.

This study seeks to use the new CERT[®] Insider Threat Lab to better understand the threat of malicious insiders. We plan on using the Insider Threat Lab for ongoing analysis and development of insider threat controls. Part of this ongoing effort involves detailed study of the types of crimes cataloged by CERT, as well as extending previous work in behavioral modeling efforts to better explain how insiders behave and carry out their attacks on organizations. This initial study analyzes the current trends of how insiders actually steal IP.

We analyze 50 incidents involving insiders who stole IP to better understand the trends and methods insiders use to exfiltrate sensitive data from an organization. This work shows how this technical data set, at multiple points, supports the behavioral findings by Moore and associates in *Insider Theft of Intellectual Property for Business Advantage: A Preliminary Model*,¹ which developed initial models and analysis of insider theft of IP cases. This research expands upon those initial findings by studying the business assets, types of IP that have been targeted, and exfiltration methods.

These findings link back to best practices and classes of controls that could mitigate the risk of these types of incidents. The following are some key observations from this work.

Observation 1: Methods used by malicious insiders to steal IP ranged widely. In the 50 cases studied, the top three methods that insiders used to steal sensitive data were

- email from work: 30 percent
- removable media: 30 percent
- remote network access: 28 percent

Observation 2: Insider use of both personal and work email remains a primary method for using networked resources to quickly exfiltrate information from an organization.

Observation 3: Of all the cases of theft of IP in our sample, 28 percent involved remote network access.

[®] CERT is a registered mark owned by Carnegie Mellon University.

¹ <http://sunsite.informatik.rwth-aachen.de/Publications/CEUR-WS/Vol-469/paper1.pdf>

Abstract

Since 2001, the Insider Threat team at the Software Engineering Institute's CERT[®] program has built an extensive library and comprehensive database containing more than 550 cases of insider crimes. More than 80 of those crimes involved theft of an organization's intellectual property by a malicious insider. These crimes can be particularly damaging to an organization because it is often difficult or impossible to recover from a loss of confidentiality. This report provides an overview of techniques employed by malicious insiders to steal intellectual property, including the types of assets targeted and the methods used to remove the information from a victim organization's control. The report closes with a brief discussion of mitigating factors and strategic items that an organization should consider when defending against insider attacks on intellectual property.

1 Introduction

1.1 Background

In 2001 the CERT[®] Program, part of Carnegie Mellon University's Software Engineering Institute, conducted a joint study on insider threats with the United States Secret Service. This original study resulted in the collection and analysis of close to 150 cases. Since 2001 the number of insider threat cases collected by CERT researchers has increased to well beyond 550. In its nearly 10 years of studying the insider threat problem, CERT has published several behavioral models for the various types of crimes the CERT[®] Insider Threat team tracks, published best practices for preventing insider attacks, and developed a suite of public services ranging from insider threat assessments to instructional workshops. The majority of this work, however, has not addressed the granular technical issues associated with insider crimes. Instead, as in the case of systems dynamics models, we have looked at patterns of behavior by both the insiders and their victim organizations that ultimately lead to the commission of a crime.

We turn an eye toward the operational staff, and its leadership, who are tasked with defending against insider threats. Recognizing that behavioral models alone are not the whole solution, we intend this work to expose detailed information about the technical threats and vulnerabilities exploited by insiders in a sample of cases. We also briefly survey some of the currently available tools that claim to have insider threat defensive capabilities and discuss where they do or do not align with our data.

We seek to use the new CERT[®] Insider Threat Lab to better understand the threat of malicious insiders. We plan on using the Insider Threat Lab for ongoing analysis and development of insider threat controls. Part of this ongoing effort involves detailed study of the types of crimes cataloged by CERT. It also involves extending previous work in behavioral modeling efforts to better explain how insiders behave and carry out their attacks on organizations. This report details one of this effort's initial studies, which analyzed the current trends of how insiders actually steal IP.

1.2 Definitions

One of the primary difficulties in studying insider threats seems to be reaching agreement on whom or what should be considered an insider at an organization. Some previously considered issues include whether contract employees are insiders, whether malware is a form of insider threat (decoupled from the human who wrote or planted it), and whether an outsider who has infiltrated an organization's network becomes an insider after establishing a foothold on the victim's network. Not discounting these items, CERT has defined insider threat as the following: "A malicious insider is a current or former employee, contractor, or other business partner who has or had authorized access to an organization's network, system or data and intentionally exceeded or misused that access in a manner that negatively affected the confidentiality, integrity, or availability of the organization's information or information systems" [Cappelli 2009].

[®] CERT is a registered mark owned by Carnegie Mellon University.

While we recognize that this definition excludes some areas previously expressed as questionable cases of insider threat, we find the definition to be a strong starting point for all of our analysis given its unambiguous boundaries. All cases in our Insider Threat database meet this definition, and references to the term “insider” throughout the remainder of this report should be considered to use this definition.

CERT also defines various categories of insider threats based on the outcome, intended or actual, of an insider’s attack. We find that these categories are important for analysis. For instance, saboteurs and thieves behave in some fundamentally different ways and use different tools and techniques. The three core categories of study are IT sabotage, fraud, and theft of intellectual property (IP). The Insider Threat team also catalogs cases that do not fit any of the three core case types, or lack sufficient information to be categorized, under a “miscellaneous” category. As with all cases in the database, we regularly update these cases as new information becomes available and categorize them as appropriate. The Insider Threat team also studies cases of national security espionage involving classified information. These last cases, however, fall outside the scope of this effort. Table 1 defines each type of crime.

Table 1: Types of Insider Crimes

<u>Type of Crime</u>	<u>Description</u>
IT Sabotage	An insider’s use of IT to direct specific harm at an organization or an individual.
Insider Fraud	An insider’s use of IT for the unauthorized modification, addition, or deletion of an organization’s data (not programs or systems) for personal gain, or theft of information that leads to fraud (identity theft, credit card fraud).
Theft of IP	An insider’s use of IT to steal IP from the organization. This category includes industrial espionage involving insiders.
Miscellaneous	Cases that do not fit well into the three primary categories.

1.3 CERT® Insider Threat Database

The CERT Insider Threat team tracks cases of insider threat in each of the three core areas, as well as other cases, in an Insider Threat database maintained by CERT staff. This database recently grew to more than 550 cases of insider crime. The Insider Threat team continuously catalogs new cases and updates old cases to ensure we are performing analysis on the most current and accurate data set possible. The database also tracks a fairly large set of attributes for each cataloged case added to the repository. At a high level, those attributes include granular details from within several areas of interest including, but not limited to

- details of the insider’s behavior and interactions with coworkers
- vulnerabilities in organization systems that the insider was able to exploit in the attack
- unmet expectations related to job conditions
- technical and nontechnical methods used by the insider
- evidence of planning and deception
- how the incident was detected or reported
- types of information or assets that were stolen or targeted
- information about coconspirators and recruitment

It should be noted that while we can collect a wide array of information about insider cases, it is not always possible to collect all the relevant details of an insider case. The source material we

use to catalog cases ranges from detailed investigative notes and interviews to media reports, the latter often lacking some of the granular details about the execution of an attack. That said, the database does have a mechanism for computing a rough quality score so that analysts can focus on improving cases that have the least information or lowest quality of information.

Figure 1 shows the breakdown of cases by category as of the time of this report's publication. While fraud constitutes the largest category of case material, it is one of the least technically interesting categories. These cases often involve nontechnical, low-level clerical and data-entry positions that have access to confidential or sensitive organization information. Sabotage is the second largest category. These cases involve the more technically savvy insiders who frequently have privileged access to multiple essential systems at the victim organization. While the techniques employed by saboteurs were outside the scope of this study, we intend to take a closer look at them in the future. Finally, theft of IP trails as the smaller category. However, CERT has recently emphasized this area. These cases typically involve scientists and engineers with access to privileged technical information and trade secrets, or salespersons with access to critical business plans and customer information. Each group stands to profit by stealing information for the benefit of a future employer, personal financial gain, or the benefit of a foreign entity. We examine a subset of these cases in Section 2 of this paper.




Figure 1: Number of Cases in the CERT Insider Threat Databases by High-Level Category (Excluding National Security Espionage Cases)

CERT staff members are frequently asked what types of victim organizations are being attacked by insiders. The Insider Threat database can associate a case with a critical infrastructure sector, making this analysis easier to perform. Figure 2 shows the type of crime for each sector.




Figure 2: Cases in Three Major Crime Types by Sector

1.4 Considerations

As previously mentioned, it can be difficult to collect complete information for each insider threat case. There are other considerations worth bearing in mind as we proceed with this discussion.

First, while the findings we present in this report are associated with a subset of theft of IP cases, the findings do not necessarily represent all insider threat cases or even all insider theft of IP cases. Our past behavioral models clearly indicate that the types of crimes differ from each other in material ways. Further, the cases we catalog are cases that have been publicly reported. Much like security incidents at large, insider cases are underreported. In the most recent *Cybersecurity Watch Survey* conducted jointly by Deloitte, the United States Secret Service, *CSO Magazine*, and CERT, respondents reported that 72 percent of events perpetrated by insiders were handled internally without any legal action or law enforcement involvement [CSO 2010]. These cases usually do not surface to public view, so we do not have the opportunity to catalog them. While we clearly are not cataloging all cases of insider crime, we find the database to be the strongest available source of information on actual insider threat incidents.

Second, this report does not necessarily advocate any particular use of the analysis of technical observables in the sample cases. We encourage organizations considering creating technical controls, alerts, or rule sets based on this information to read *Deriving Candidate Technical Controls and Indicators of Insider Attack from Socio-Technical Models and Data* [Hanley 2011], another recent report from CERT describing a method for aligning technical descriptors with findings from behavioral models to create stronger control sets. CERT is actively developing candidate indicators, controls, and configurations based on data from analysis such as this and prior behavioral research to create informed control sets that can be customized by organizations based on their own unique environment, operational needs, and intelligence.

Organizations considering the observations in this report should consider that the collection mechanisms required to make some of these observations may have privacy, legal, and human resources implications. Any operationalization of these findings or development of control sets should include a review by appropriate management, legal, human resources, and privacy staff to ensure consistency and compliance with applicable laws and regulations.

Monitoring techniques are not a guarantee. Many of the observations we cite in this report would have indicated that an attack was either in progress or had already been completed. These observations may be most useful for triggering a response capability where resources are limited.

2 Analysis

2.1 Case Pool

The case pool for this study consisted of a sample of cases from the theft of IP category within the CERT Insider Threat database.

At the time the sample was created, the Insider Threat database contained 64 theft of IP cases. We extracted for analysis 50 of these cases whose material events of the crime would likely have created technical observables for host or network monitoring solutions to detect. Throughout our review of the case material, we maintained a constant focus on how an organization might have prevented or detected the attack, intervened to mitigate its consequences, or accelerated the recovery process.

2.2 Nontechnical Findings of Interest

In the 50 extracted cases, 94 percent of the insiders were male. The most prevalent positions held by insiders were engineer or scientist (44 percent), manager (18 percent), salesperson (14 percent), and programmer (10 percent). Most insiders who stole IP were male employees in technical positions within the victim organization. These findings roughly align with findings from the most recent theft of IP modeling research from CERT, which found that 91 percent of such insiders were male and 55 percent held technical positions [Moore 2009].

As previously stated, each case in the Insider Threat database is associated with the appropriate critical infrastructure to which the victim organization belongs. While the data shows that most sectors experienced at least one instance of insider theft of IP, the sample we used associates 54 percent of its cases with the IT and telecommunications sector. Only 8 percent of cases originated in the manufacturing sector.

In almost half of the cases in the case pool, the insider attacked during normal business hours, while in 20 percent the insider attacked outside of normal business hours. For the remaining 36 percent of the cases, it was unknown when the insider attacked. In three cases, the insider attacked both during and outside of normal working hours. Furthermore, 62 percent of the incidents occurred on-site, while 22 percent occurred through remote access. In 24 percent of the cases, the location was unknown; four cases were carried out both on-site and remotely. Most of the incidents occurred during working hours and on the organization's site.

2.3 Data Structure

The selected cases had to be transformed into a form suitable for an analysis of technical observables. We defined a basic table of attributes to extract from all cases in the sample, both for this work and other ongoing work at CERT. The case table is simple, comprises the relevant technical details of a crime, and streamlines the analysis process. The case table also allowed us to use standard formatting and abstract technical details in a uniform way across cases. For each field, we included a standard set of responses based on observations in the data. Applying this table to

an alternate data set could yield additional acceptable responses. Table 2 shows an example of a completed case table.

Table 2: Sample Case Table

Name	<Insider Name>
Incident_ID	<Integer value>
Summary	Subject failed to receive raise and a request for transfer was rejected. Subject submitted resignation and downloaded proprietary information from organization for potential use in new company. Subject used file transfer protocol (FTP) to transmit data to home computer.
Assets Attacked/Targeted	Internal business information
Source Type	Electronic documents
Method of Exfiltration	File transfer
Exfiltration Comments	Insider was able to open an FTP connection offsite to move the data outside the network.
Controls	
Prevention	Restrict access after employee resignation; clarify ownership of IP.
Detection	Monitor behavior between resignation and termination; monitor user network activity/downloads.
Response	Audit user activity; notify new organization of behavior.
Solutions	
COTS	Remote access file transfer monitoring

2.4 Data Analysis by Field

The following sections describe the process used to complete the case tables of all 50 sampled cases and present findings from each field, where applicable.

2.4.1 Case Name and ID

These items are for internal tracking only. CERT does not disclose the names of insiders or their victim organizations. The “Incident_ID” is a unique integer value assigned to each unique case.

2.4.2 Case Summary

Perhaps one of the more important fields in the case table, the case summary encapsulates the events of a crime and provides context for an analyst, who must move quickly through a large volume of cases and their technical details. The case table uses a significantly shorter summary of case events than that which appears in the actual Insider Threat database, which can be several pages long in very complex cases. The summary in the case table provides a relatively small amount of context and information about the exfiltration event itself.

2.4.3 Assets Attacked/Targeted

To enable cross-case tracking of certain types of assets that may be targeted more frequently based on ease of access, perceived value, or association with job role, the case table includes the types of information insiders steal. Because granularity of information on what an insider targeted or stole varies from case to case, we attempted to abstract all of the targets to roughly the same level by using the following categories:

- customer information
- source code

- business plans
- trade secrets
- internal organization information
- proprietary software

A code book defines each category to ensure information is abstracted uniformly from analyst to analyst and from case to case. The case table may include more than one category if the insider stole or targeted more than one type of asset. Table 3 shows the definitions.

Table 3: Assets Attacked/Targeted Responses

Type	Information
Customer Information	Includes various data about an organization's customers. For example, lists of the organization's customers, customer quotes, customer orders, and any other data relating to the organization's customers and clients. When analyzing cases, if at least one of these items were targeted, the case should include this category.
Source Code	Includes code written by the insider and source code on the network written by other programmers. Regardless of who authored the code, if the insider stole or targeted this information, the case should be included in this category.
Business Plans	Includes many forward-looking and strategic initiatives of the organization. For example, marketing plans, sales plans, and other business plans were prime targets for insiders with access to them. Cases that involved these types of documents and plans should be included in this category.
Trade Secrets	Consists of a wide range of proprietary information. For instance, an organization's product or service information, product designs, specifications, formulae, and other proprietary information should result in a case being included in this category.
Internal Organization Information	Includes any information used by organizations throughout their normal course of business activity. This can include billing information, price lists, sales brochures, and other internal information. Cases targeting this information should be categorized here.
Proprietary Software	Complete products, or components of products, written specifically for the organization that provided a business advantage over competitors. In contrast to source code, proprietary software is defined as the actual compiled and useable software. These cases were cataloged in this category if they included theft of these applications and conveying them to a competitor or other third party.

We found that by far the largest category of stolen information was trade secrets (see Figure 3). Of the 50 cases, 52 percent involved stolen trade secrets. This seems intuitive because trade secrets generally provide a company with a competitive advantage in their market. At a more granular level, it appears that insiders often target product designs, formulas, or other proprietary information within this category.

The next largest targeted category is internal organization information, such as strategies, billing information, price lists, sales brochures, and other internal organization resources. Of the 50 cases, 30 percent involved stolen internal business information. This is a generalized category that applies to many nongovernment organizations, so it is not surprising that a large percentage of cases involved this type of asset.

Insiders targeted another kind of asset, source code, in 20 percent of the cases. Confusion over source code ownership was a common theme; insiders often claimed they felt they owned the code they had written and therefore were entitled to take it with them when leaving the firm. While we do not have data on how the victim organizations enforced ownership of their source

code as IP (if at all), it is critical to consider how organizations can clearly communicate who owns products developed by employees for the organization.

Insiders targeted proprietary software, customer information, and business plans in 14 percent, 12 percent, and 6 percent of the cases, respectively. Customer information included customer lists, orders, and other data. The fact that customer information was targeted in only 12 percent of cases may initially seem surprising. However, we did not include customer information that includes personally identifiable information (PII), such as high-value items like Social Security numbers or card verification value (CVV) numbers on credit cards, in theft of IP cases. Those are frequently coded as fraud because identity theft is usually the end goal. This may explain the lower numbers for stolen customer information.




Figure 3: Assets Attacked by Insiders

2.4.4 Methods of Exfiltration

Each case was also examined to determine the primary method of exfiltrating data from within the organization's boundary. Table 4 shows the definitions of each method. We identified six major methods of exfiltration from the data sample.

- email
- removable media
- printed documents
- remote network access
- file transfer
- laptops

Table 4: Exfiltration Methods

Type	Information
Email	The insider attempted to exfiltrate information through the insider's work email account. The email may have been sent to a personal email account or directly to a competitor. Insiders could have used email attachments or the body of the message to transmit the sensitive information out of the network.
Removable Media	Common removable media types were categorized as thumb drives, CDs, and removable hard drives. If the exact type of removable media was not known, the removable media category was still used but was identified as "unknown."
Printed Documents	Whether an insider printed a document or printed a screenshot of sensitive information, these cases involved insiders physically removing IP from the victim organization.
Remote Network Access	The insider remotely accessed the network through a virtual private network (VPN) or other remote channel to download sensitive information from an off-site location. The key aspect of this method was that an insider must have transferred data out of the network from a remote location.
File Transfer	The insider transferred sensitive information out of the network using the web, file transfer protocol (FTP), or other methods. The difference between remote network access and file transfer is that file transfer involves an insider who was at work, was on the company network, and transferred a file outside of the network. Although email could potentially fit into this category, we felt that email could be considered a separate issue and category.
Laptops	The insider attempted to exfiltrate data by downloading IP onto a laptop at work and bringing it outside the workplace. For example, one insider was developing an application for the company on a laptop and later purposefully leaked that code. In other cases, the insiders simply downloaded sensitive files onto their laptops for personal or business use later.

We further condensed these categories to see whether exfiltration largely involved the host, the network, or physical removal by the insider. Figure 4 shows the totals.




Figure 4: Exfiltration Methods (Aggregate)

2.4.4.1 Network Data Exfiltration

Data exfiltration over the network was the most common method of removing information from an organization. Of the cases we analyzed, 54 percent involved removal of data from the organization using email, a remote network access channel (originating externally), or a network file

transfer (originating inside the network). In the case of email, 15 of the insiders used this method to move data off the network. This took one of two forms: an insider emailing IP from the insider's work email account to a personal email account for later use, or an insider directly emailing IP to a competitor from the insider's work email account.

In cases where insiders used personal accounts, the insider first emailed sensitive information to a personal email account from the insider's work email account. Then the insider would transfer it to a competitor or use it to start the insider's own business. Some insiders used the sensitive information to gain a better position at a competing organization, while others emailed the IP to their new employer to give it an unfair business advantage. For example, an insider in one case sent customer lists and source code he had written from his work email account to his personal email account. During this time, the insider was recruited by a competing organization for employment. The insider accepted the offer and took the customer lists and source code to his new employer to help him succeed at his new job.

In the second type of case involving email, the insiders simply emailed sensitive documents directly to a competitor from their work email account. In one case, an insider asked his superiors for confidential data about his present company's product costs and materials. Two months later, he accepted a new job with a competitor. The original employer warned the insider against taking or distributing any of its proprietary information. However, the insider emailed internal business information from his old work email account to two of his new supervisors before he started at the new company.

Interestingly, 47 percent of cases involving exfiltration via email also involved an additional type of data exfiltration, perhaps suggesting that an analyst or operator who suspects an insider is stealing from the organization should check other communication channels for similar activity. Most frequently, the additional exfiltration path involved stealing information on a laptop, but use of remote access channels and theft of printed documents each appeared twice in combination with theft via email.

The second most frequent network exfiltration method was remote network access. In 14 cases, the insider remotely accessed the network to exfiltrate information. Many of these cases occurred immediately before resignation or shortly after acceptance of a new job at a competitor. In 36 percent of these cases, the remote connections were established after normal working hours; in 29 percent of these cases, the time of exfiltration was unknown. During the remote sessions, users were able to download sensitive documents to their remote computers. In one case, an insider and a coworker were employed as contract software developers for the victim organization. After their contracts ended, the victim organization failed to terminate their remote access to the network. Both insiders claimed that certain programs they developed belonged to them, and they suddenly requested that the organization cease using them. The company continued to use the applications, and the insider and coworker remotely accessed and downloaded the proprietary source code they claimed to own.

The least common method of network data exfiltration was transferring data outside the network through outbound channels such as FTP, the web, or instant messaging. Only three of these types of incidents existed in the case pool, and they were all perpetrated by more technically skilled insiders. In the first case, an insider worked at an investment banking organization as a computer programmer. Prior to the incident, the insider had submitted his letter of resignation to his manag-

er. The insider used a script that copied, compressed, and merged files containing source code and then encrypted, renamed, and uploaded the files using FTP to an external file hosting server. The second case involved an insider who transferred trade secrets and source code to a password-protected website using standard HTTP. The insider intended to start a side business with the company's stolen IP. In the third case, the insider failed to receive a raise, and his request for transfer was rejected. The insider submitted his resignation and downloaded proprietary information from his organization for potential use in a new job. He used FTP to transfer the data to his home computer.

2.4.4.2 Host Data Exfiltration

Host-based exfiltration was the second most common method for removing sensitive data from the company. Of the 50 cases, 25 incidents involved an insider removing data from a host computer and leaving the company with it.

In these cases, insiders often used their laptops to remove data from the organization. We had difficulty determining the exact ownership and authorization of the laptops used. However, 10 of the insiders used laptops taken from the organization's site during normal working hours. Five insiders transferred proprietary software and source code, and five insiders removed sensitive documents from the organization.

In one case, the insider worked for a consulting company and stole proprietary software programs by downloading them to a laptop. The insider attempted to disguise the theft by deleting references to the victim organization contained in the program and then attempted to sell portions of the program to a third party for a large sum of money.

Another case involved an insider who accessed and downloaded trade secrets to his laptop after he accepted an offer from a foreign competitor. He did not notify his current organization of his new job until two weeks before termination. He continued to steal information until he left.

By far, the most common method of host-based exfiltration in the case pool was removable media. Of the cases identified, 80 percent involved trade secrets. Of the insiders who stole trade secrets using removable media, 67 percent took the stolen trade secrets to a competitor. The type of removable media used varied. Where information was available, we determined that insiders most often used writable CDs. Thumb drives and external hard disks were used in four cases each.

The type of removable media insiders have used to copy and steal IP has changed over time. Insiders primarily used CDs prior to 2005. Since 2005, however, most insiders using removable media to steal IP use thumb drives and external hard drives. This trend indicates that changes in technology are providing new and easier methods of stealing data from host computers.

In one case, an insider resigned from his organization after accepting a position at another organization. The insider downloaded personal files as well as the organization's proprietary information onto CDs. Despite signing a nondisclosure agreement, the insider took the trade secrets to a competitor.

In a similar example, an insider received an offer from a competitor three months prior to resignation. The insider lied about his new position and employment status to coworkers. Only days before leaving the organization, the insider convinced a coworker to download the insider's files to

an external hard drive, supposedly to free up disk space. The insider came into work at unusual hours to download additional proprietary information onto a CD. Finally, the insider took this information with him to his new position at a competing organization.

2.4.4.3 Physical Exfiltration

We found 4 of the 50 cases to have involved some sort of physical exfiltration of sensitive IP. This number is small, partially because our case selection criteria filtered for cases that emphasized technical observables. We found that physical exfiltration usually occurs in conjunction with some other form of exfiltration that would have produced a more obvious network or host-based observable event.

2.5 Analysis Across Multiple Fields

Once we determined what kinds of assets were stolen and how, we determined what methods of exfiltration were associated with the different asset types. Figure 5 shows the results.




Figure 5: Exfiltration by Asset Type

Several interesting findings surfaced. In particular, business plans were stolen almost exclusively through network methods, particularly remote access from off-site. Conversely, proprietary software and source code involve a much higher use of non-network methods. This may be due in part to the volume of data associated with different asset types. Software and source code files are often large, but business plans are usually smaller documents that are easier to move over a VPN or as an email attachment. Enumerating the most frequent methods by which particular assets are exfiltrated may help steer monitoring strategies with respect to computers that house particular types of assets or are allowed to access given assets over the network.

2.5.1 Concealment

Some evidence in the data suggests that some insiders attempted to conceal their theft of sensitive information through various actions, though this evidence is not as strong as that of other observations in this study. These cases signify a clear intent to operate covertly, implying the insiders may have known their actions were wrong.

Of the 50 cases analyzed, 6 contained clear information about concealment methods used by the insider. In one case, an insider was arrested by federal authorities after stealing product design documents and transferring them to a foreign company where he was to be employed. After being arrested, the insider instructed a friend to log into the insider's personal email account, which was used in the exfiltration, and delete hundreds of emails related to the incident. Another case involved an insider who used an encryption suite to mask the data he had stolen when moving it off the network.

3 Considerations for Mitigation

3.1 Recommendations

While it was not the intent of this study to suggest detailed technical controls, some trends identified in the data are clearly worth discussion. Any organization considering implementing controls based upon these findings should develop insider threat strategies involving more than technical tools alone. Such controls should be endorsed by senior leadership in HR, legal, physical security, and other relevant areas of the organization. The findings of this study may not lead directly to configurations and controls capable of preventing insider crime, but they may still be useful for incident responders when reconstructing the events of an incident or directing their efforts based on past cases. The CERT *Common Sense Guide* [Cappelli 2009] provides general guidance for insider threat defensive strategies.

One of the most critical findings of this study is that despite improvements in communication, desktop computing, and mobile computing solutions, many insiders continue to steal information using removable media. It is unlikely that the victim organizations prohibited removable media in their daily computing environments. Organizations should consider, at a minimum, having some measure of employee use of removable media. Understanding who requires removable media and for what purposes can help an organization determine what may constitute normal and healthy business use. Inventory control, as it pertains to removable media, may also be helpful. For example, an organization could allow use of removable media only on company-owned devices prohibited from leaving the facility. Organizations requiring the highest-assurance environment should consider disallowing removable media, or allowing it only in special situations that are carefully audited.

Most cases that involved use of the network to perpetrate the theft involved email and remote access over VPN. Given that several cases involved email to a direct competitor, firms should consider at least tracking, if not blocking, email to and from competing organizations. Our cases did not explicitly show sophisticated concealment methods, such as use of proxies or extensive use of personal, web-based email services. However, we did find that insiders periodically leverage their personal, web-based email as an exfiltration method. Employers should carefully consider the balance between security and personal use of email and web services from organization systems.

According to the theft of IP models created by CERT, most insiders steal IP within 30 days of leaving an organization. Organizations should consider a more targeted monitoring strategy for users who have already given notice of their exit. Further, organizations should consider inspecting available log traffic for any indicators of suspicious access, large file transfers, suspicious email traffic, after-hours access, or use of removable media. Central logging appliances and event correlation engines may help craft automated queries that reduce an analyst's workload for routinely inspecting this data.

Finally, organizations should consider a review of access termination policies associated with employee exit procedures. Several cases provided evidence that insiders remotely accessed systems by using previously authorized accounts that were not terminated upon the employee's exit. Pre-

cautions against this kind of incident would seem to be common sense, but this trend continues to manifest in newly cataloged cases.

4 Conclusion and Future Work

This study analyzed several interesting technical observations from insider threat cases of theft of IP. The results provided some insight into insiders' tactics for stealing information from organizations. We have also briefly presented potential strategies, tailored to the data of this study, that an organization might consider when implementing monitoring strategies or pursuing its own independent insider threat control development.

In the coming months, CERT will be using a recently built Insider Threat Lab to perform live testing of insider threat tools and tool configurations against re-creations of actual insider events from our database. This effort will develop more informed technical control sets organizations can implement to improve their security posture. These technical controls will be derived in part from behavioral models and in part from the descriptors in the case material. They may be of significant benefit to organizations seeking to combat insider threats with limited resources. We intend to pilot several of these control sets on live organization networks as well to provide additional validation and implementation guidance for organizations seeking assistance in this arena.

References

URLs are valid as of the publication date of this document.

[Cappelli 2009]

Cappelli, Dawn; Moore, Andrew; Trzeciak, Randall; & Shimeall, Timothy J. *Common Sense Guide to Prevention and Detection of Insider Threats, 3rd Edition – Version 3.1*. Software Engineering Institute, Carnegie Mellon University, 2009. <http://www.cert.org/archive/pdf/CSG-V3.pdf>

[CSO 2010]

CSO Magazine; Software Engineering Institute; Deloitte; & U.S. Secret Service. *2010 Cybersecurity Watch Survey: Cybercrime Increasing Faster than Some Company Defenses*. CXO Media Inc., 2010. <http://www.cert.org/archive/pdf/ecrimesummary10.pdf>

[Hanley 2011]

Hanley, Michael. *Deriving Candidate Technical Controls and Indicators of Insider Attack from Socio-Technical Models and Data* (CMU/SEI-2011-TN-003). Software Engineering Institute, Carnegie Mellon University, 2011. <http://www.sei.cmu.edu/library/abstracts/reports/11tn003.cfm>

[Moore 2009]

Moore, Andrew P.; Cappelli, Dawn M.; Caron, Thomas C.; Shaw, Eric; & Trzeciak, Randall F. “Insider Theft of Intellectual Property for Business Advantage: A Preliminary Model” 1-21. *Proceedings of the 1st International Workshop on Managing Insider Security Threats* (MIST 2009). Purdue University, West Lafayette, IN, June 2009. CEUR-WS Proceedings vol. 469, 2009. <http://sunsite.informatik.rwth-aachen.de/Publications/CEUR-WS/Vol-469/paper1.pdf>

REPORT DOCUMENTATION PAGE			<i>Form Approved</i> <i>OMB No. 0704-0188</i>	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.				
1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE February 2011		3. REPORT TYPE AND DATES COVERED Final
4. TITLE AND SUBTITLE An Analysis of Technical Observations in Insider Theft of Intellectual Property Cases			5. FUNDING NUMBERS FA8721-05-C-0003	
6. AUTHOR(S) Michael Hanley, Tyler Dean, Will Schroeder, Matt Houy, Randall F. Trzeciak, Joji Montelibano				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Software Engineering Institute Carnegie Mellon University Pittsburgh, PA 15213			8. PERFORMING ORGANIZATION REPORT NUMBER CMU/SEI-2011-TN-006	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) HQ ESC/XPK 5 Eglin Street Hanscom AFB, MA 01731-2116			10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12A DISTRIBUTION/AVAILABILITY STATEMENT Unclassified/Unlimited, DTIC, NTIS			12B DISTRIBUTION CODE	
13. ABSTRACT (MAXIMUM 200 WORDS) Since 2001, the Insider Threat team at the Software Engineering Institute's CERT® program has built an extensive library and comprehensive database containing more than 550 cases of insider crimes. More than 80 of those crimes involved theft of an organization's intellectual property by a malicious insider. These crimes can be particularly damaging to an organization because it is often difficult or impossible to recover from a loss of confidentiality. This report provides an overview of techniques employed by malicious insiders to steal intellectual property, including the types of assets targeted and the methods used to remove the information from a victim organization's control. The report closes with a brief discussion of mitigating factors and strategic items that an organization should consider when defending against insider attacks on intellectual property.				
14. SUBJECT TERMS insider threat, information security, system dynamics, behavioral modeling, security controls, counterintelligence, security metrics			15. NUMBER OF PAGES 35	
16. PRICE CODE				
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL	